

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 December 2001 (06.12.2001)

PCT

(10) International Publication Number
WO 01/93028 A2

(51) International Patent Classification⁷: G06F 9/38

(21) International Application Number: PCT/US01/17645

(22) International Filing Date: 30 May 2001 (30.05.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/208,428 31 May 2000 (31.05.2000) US
09/761,226 16 January 2001 (16.01.2001) US

(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901
San Antonio Road, Palo Alto, CA 94303 (US).

(72) Inventors: CHAUDHRY, Shailender; 1200 Gough St.,
Apt. 10F, San Francisco, CA 94109 (US). TREMBLAY,
Marc; 140 Hanna Way, Menlo Park, CA 94025 (US).

(74) Agent: PARK, Richard; 508 2nd Street, Suite 201, Davis,
CA 95616 (US).

(81) Designated States (*national*): AE, AG, AI, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ,
DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR,
HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR,
LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ,
NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM,
TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

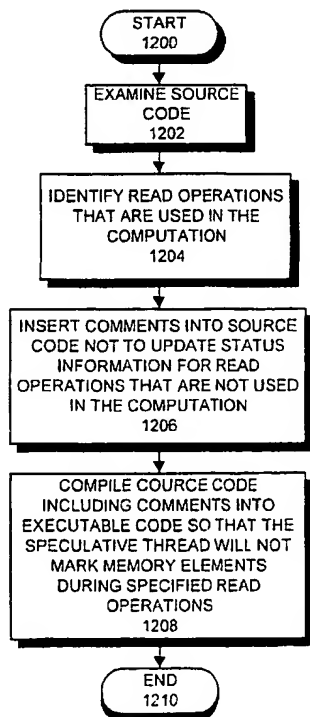
(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian
patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European
patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE,
IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF,
CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: MARKING MEMORY ELEMENTS BASED UPON USAGE OF ACCESSED INFORMATION DURING SPECULA-
TIVE EXECUTION



(57) Abstract: One embodiment of the present invention provides a system that marks memory elements based upon how information retrieved from the memory elements affects speculative program execution. This system operates by allowing a programmer to examine source code that is to be compiled into executable code for a head thread that executes program instructions, and for a speculative thread that executes program instructions in advance of the head thread. During read operations to memory elements by the speculative thread, this executable code generally causes the speculative thread to update status information associated with the memory elements to indicate that the memory elements have been read by the speculative thread. Next, the system allows the programmer to identify a given read operation directed to a given memory element, wherein a given value retrieved from the given memory element during the given read operation does not affect subsequent execution of the speculative thread. The programmer is then allowed to insert a hint into the source code specifying that the speculative thread is not to update status information during the given read operation directed to the given memory element. Next, the system compiles the source code, including the hint, into the executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

WO 01/93028 A2

MARKING MEMORY ELEMENTS BASED UPON USAGE OF ACCESSED INFORMATION DURING SPECULATIVE EXECUTION

5

BACKGROUND

Field of the Invention

10 The present invention relates to compilers and techniques for improving computer system performance. More specifically, the present invention relates to a method and an apparatus that supports marking a memory element based upon how information retrieved from the object is used in order to facilitate speculative program execution.

15

Related Art

As increasing semiconductor integration densities allow more transistors to be integrated onto a microprocessor chip, computer designers are investigating different methods of using these transistors to increase computer system performance. Some recent computer architectures exploit "instruction level parallelism," in which a single central processing unit (CPU) issues multiple instructions in a single cycle. Given proper compiler support, instruction level parallelism has proven effective at increasing computational performance across a wide range of computational tasks. However, inter-instruction dependencies generally limit the performance gains realized from using instruction level parallelism to a factor of two or three.

25

Another method for increasing computational speed is "speculative execution" in which a processor executes multiple branch paths simultaneously, or predicts a branch, so that the processor can continue executing without waiting for the result of the branch operation. By reducing dependencies on branch conditions, speculative execution can increase the total number of instructions issued.

30

Unfortunately, conventional speculative execution typically provides a limited performance improvement because only a small number of instructions can be speculatively executed. One reason for this limitation is that conventional speculative execution is typically performed at the basic block level, and basic blocks tend to include only a small number of instructions. Another reason is that conventional hardware structures used to perform speculative execution can only accommodate a small number of speculative instructions.

What is needed is a method and apparatus that facilitates speculative execution of program instructions at a higher level of granularity so that many more instructions can be speculatively executed.

One challenge in designing a system that supports speculative execution is to detect a rollback condition. A rollback condition can occur in a number of situations. For example, a rollback condition can occur when a speculative thread that is executing program instructions in advance of a head thread reads from a memory element before the head thread performs a store to the memory element. In this case, the speculative thread must "rollback" so that it can read the value stored by the head thread. A rollback condition can be detected by "marking" memory elements as they are read by the speculative thread so that the head thread can subsequently determine if the memory elements have been read by the speculative thread.

However, not all data values that are read by a speculative thread affect subsequent execution of the speculative thread. For example, a lookup into a hash table by the speculative thread may involve a number of preliminary accesses to memory elements in the hash table before a desired hash table entry is located. Note that these preliminary accesses do not affect subsequent actions of the speculative thread. Hence, marking these memory elements can cause future accesses to these memory elements to generate unnecessary rollbacks, which can reduce the performance gains arising from speculative execution.

What is needed is a method and an apparatus for selectively marking accessed memory elements based upon whether information retrieved from the accessed memory elements affects subsequent actions of the speculative thread.

3

SUMMARY

One embodiment of the present invention provides a system that marks memory elements based upon how information retrieved from the memory elements affects speculative program execution. This system operates by allowing a programmer to examine source code that is to be compiled into executable code for a head thread that executes program instructions, and for a speculative thread that executes program instructions in advance of the head thread. During read operations to memory elements by the speculative thread, this executable code generally causes the speculative thread to update status information associated with the memory elements to indicate that the memory elements have been read by the speculative thread. Next, the system allows the programmer to identify a given read operation directed to a given memory element, wherein a given value retrieved from the given memory element during the given read operation does not affect subsequent execution of the speculative thread. The programmer is then allowed to insert a hint into the source code specifying that the speculative thread is not to update status information during the given read operation directed to the given memory element. Next, the system compiles the source code, including the hint, into the executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

In one embodiment of the present invention, during a subsequent write operation to a memory element by the head thread, the executable code causes the head thread to check status information associated with the memory element to determine if the memory element has been read by the speculative thread. If the memory element has been read by the speculative thread, the system causes the speculative thread to roll back in order to read a result of the write operation.

In one embodiment of the present invention, inserting the hint into the source code involves inserting a comment into the source code that causes a compiler to generate executable code for the speculative thread that does not to mark the given memory element during the given read operation.

In one embodiment of the present invention, compiling the source code into the executable code involves setting a look attribute associated with the given memory element to indicate that the speculative thread is not to mark the given memory element. Note that the code for marking (extra stores and loads) is not in the executable code. Hence, it results in increased speed of execution of the speculative thread.

In one embodiment of the present invention, the executable code causes a join operation between the head thread and the speculative thread when the head thread reaches a point in a program where the speculative thread began executing. This join operation causes state associated with the speculative thread to be merged with state associated with the head thread.

In one embodiment of the present invention, the system additionally integrates either the executable code or the source code into a system library.

In one embodiment of the present invention, the memory elements include objects defined within an object-oriented programming system.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 illustrates a computer system including two central processing units sharing a common data cache in accordance with an embodiment of the present invention.

FIG. 2A illustrates sequential execution of methods by a single thread.

FIG. 2B illustrates space and time dimensional execution of a method in accordance with an embodiment of the present invention.

FIG. 3 illustrates the state of the system stack during space and time dimensional execution of a method in accordance with an embodiment of the present invention.

FIG. 4 illustrates how memory is partitioned between stack and heap in accordance with an embodiment of the present invention.

FIG. 5 illustrates the structure of a primary version and a space-time dimensioned version of an object in accordance with an embodiment of the present invention.

FIG. 6 illustrates the structure of a status word for an object in accordance with an embodiment of the present invention.

FIG. 7 is a flow chart illustrating operations involved in performing a write to a memory element by a head thread in accordance with an embodiment of the present invention.

FIG. 8 is a flow chart illustrating operations involved in performing a read to a memory element by a speculative thread in accordance with an embodiment of the present invention.

FIG. 9 is a flow chart illustrating operations involved in performing a write to a memory element by a speculative thread in accordance with an embodiment of the present invention.

FIG. 10 is a flow chart illustrating operations involved in performing a join between a head thread and a speculative thread in accordance with an embodiment of the present invention.

FIG. 11 is a flow chart illustrating operations involved in performing a join between a head thread and a speculative thread in accordance with another embodiment of the present invention.

FIG. 12 is a flow chart illustrating the process of producing code so that a speculative thread does not mark selected memory elements in accordance with an embodiment of the present invention.

FIG. 13 is a flow chart illustrating the process of compiling source code into executable code so that a speculative thread does not mark selected memory elements in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not

intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium
5 that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital video discs), and computer instruction signals embodied in a carrier wave. For example, the carrier wave may carry information across a communications network, such as the Internet.

10

Computer System

FIG. 1 illustrates a computer system including two central processing units (CPUs) 102 and 104 sharing a common data cache 106 in accordance with an embodiment of the present invention. In this embodiment, CPUs 102 and 104 and
15 data cache 106 reside on silicon die 100. Note that CPUs 102 and 104 may generally be any type of computational devices that allow multiple threads to execute concurrently. In the embodiment illustrated in FIG. 1, CPUs 102 and 104 are very long instruction word (VLIW) CPUs, which support concurrent execution of multiple instructions executing on multiple functional units. VLIW CPUs 102 and 104 include
20 instruction caches 112 and 120, respectively, containing instructions to be executed by VLIW CPUs 102 and 104.

VLIW CPUs 102 and 104 additionally include load buffers 114 and 122 as well as store buffers 116 and 124 for buffering communications with data cache 106. More specifically, VLIW CPU 102 includes load buffer 114 for buffering loads
25 received from data cache 106, and store buffer 116 for buffering stores to data cache 106. Similarly, VLIW CPU 104 includes load buffer 122 for buffering loads received from data cache 106, and store buffer 124 for buffering stores to data cache 106.

VLIW CPUs 102 and 104 are additionally coupled together by direct communication link 128, which facilitates rapid communication between VLIW
30 CPUs 102 and 104. Note that direct communication link 128 allows VLIW CPU 102

to write into communication buffer 126 within VLIW CPU 104. It also allows VLIW CPU 104 to write into communication buffer 118 within VLIW CPU 102.

In the embodiment illustrated in FIG. 1, Data cache 106 is fully dual-ported allowing concurrent read and/or write accesses from VLIW CPUs 102 and 104. This
5 dual porting eliminates cache coherence delays associated with conventional shared memory architectures that rely on coherent caches.

In one embodiment of the present invention, data cache 106 is a 16K byte 4-way set-associative data cache with 32 byte cache lines.

Data cache 106, instruction caches 112 and instruction cache 120 are coupled
10 through switch 110 to memory controller 111. Memory controller 111 is coupled to dynamic random access memory (DRAM) 108, which is located off chip. Switch 110 may include any type of circuitry for switching signal lines. In one embodiment of the present invention, switch 110 is a cross bar switch.

The present invention generally applies to any computer system that supports
15 concurrent execution by multiple threads and is not limited to the illustrated computing system. However, note that data cache 106 supports fast accesses to shared data items. These fast accesses facilitate efficient sharing of status information between VLIW CPUs 102 and 104 to keep track of accesses to versions of memory objects.

20

Space-Time Dimensional Execution of Methods

FIG. 2A illustrates sequential execution of methods in a conventional computer system by a single head thread 202. In executing a program, head thread 202 executes a number of methods in sequence, including method A 204, method B
25 206 and method C 208.

In contrast, FIG. 2B illustrates space and time dimensional execution of a method in accordance with an embodiment of the present invention. In FIG. 2B, head thread 202 first executes method A 204 and then executes method B 206. (For this example, assume that method B 206 returns a void or some other value that is not
30 used by method C 208. Alternatively, if method C 208 uses a value returned by

method B206, assume that method C 208 uses a predicted return value from method B 206.)

As head thread 202 executes method B 206, speculative thread 203 executes method C 208 in a separate space-time dimension of the heap. If head thread 202
5 successfully executes method B 206, speculative thread 203 is joined with head thread 202. This join operation involves causing state associated with the speculative thread 203 to be merged with state associated with the head thread 202 and the collapsing of the space-time dimensions of the heap.

If speculative thread 203 for some reason encounters problems in executing
10 method C 208, speculative thread 203 performs a rollback operation. This rollback operation allows speculative thread 203 to reattempt to execute method C 208. Alternatively, head thread 202 can execute method C 208 non-speculatively and speculative thread 203 can execute a subsequent method.

There are a number of reasons why speculative thread 203 may encounter
15 problems in executing method C 208. One problem occurs when head thread 202 executing method B 206 writes a value to a memory element (object) after speculative thread 203 has read the same memory element. The same memory element can be read when the two space-time dimensions of the heap are collapsed at this memory element at the time of the read by speculative thread 203. In this case, speculative
20 thread 203 should have read the value written by head thread 202, but instead has read a previous value. In this case, the system causes speculative thread 203 to roll back so that speculative thread 203 can read the value written by head thread 202.

Note that the term “memory element” generally refers to any unit of memory that can be accessed by a computer program. For example, the term “memory
25 element” may refer to a bit, a byte or a word memory, as well as a data structure or an object defined within an object-oriented programming system.

FIG. 3 illustrates the state of the system stack during space and time dimensional execution of a method in accordance with an embodiment of the present invention. Note that since programming languages such as the Java programming
30 language do not allow a method to modify the stack frame of another method, the system stack will generally be the same before method B 206 is executed as it is

before method C 208 is executed. (This is not quite true if method B 206 returns a parameter through the system stack. However, return parameters are can be explicitly dealt with as is described below.) Referring the FIG. 3, stack 300 contains method A frame 302 while method A 204 is executing. When method A 204 returns, method B 206 commences and method A frame 302 is replaced by method B frame 304. Finally, when method B 206 returns, method C 208 commences and method B frame 304 is replaced by method C frame 306. Note that since stack 300 is the same immediately before method B 206 executed as it is immediately before method C 208 is executed, it is possible to execute method C 208 using a copy of stack 300 without first executing method B 206.

In order to undo the results of speculatively executed operations, updates to memory need to be versioned. The overhead involved in versioning all updates to memory can be prohibitively expensive due to increased memory requirements, decreased cache performance and additional hardware required to perform the versioning.

Fortunately, not all updates to memory need to be versioned. For example, updates to local variables -- such as a loop counter -- on a system stack are typically only relevant to the thread that is updating the local variables. Hence, even for speculative threads versioning updates to these local variables is not necessary.

When executing programs written in conventional programming languages, such as C, it is typically not possible to determine which updates are related to the heap, and which updates are related to the system stack. These programs are typically compiled from a high-level language representation into executable code for a specific machine architecture. This compilation process typically removes distinctions between updates to heap and system stack.

The same is not true for new platform-independent computer languages, such as the JAVATM programming language distributed by SUN Microsystems, Inc. of Palo Alto, California. (Sun, the Sun logo, Sun Microsystems, and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.) A program written in the Java programming language is typically compiled into a class file containing Java byte codes. This class file can be

transmitted over a computer network to a distant computer system to be executed on the distant computer system. Java byte codes are said to be “platform-independent,” because they can be executed across a wide range of computing platforms, so long as the computing platforms provide a Java virtual machine.

5 A Java byte code can be executed on a specific computing platform by using an interpreter or a just in time (JIT) compiler to translate the Java byte code into machine code for the specific computing platform. Alternatively, a Java byte code can be executed directly on a Java byte code engine running on the specific computing platform.

10 Fortunately, a Java byte code contains more syntactic information than conventional machine code. In particular, the Java byte codes differentiate between accesses to local variables in the system stack and accesses to the system heap. Furthermore, programs written in the Java programming language do not allow conversion between primitive and reference types. Such conversion can make it hard
15 to differentiate accesses to the system stack from accesses to the system heap at compile time.

Data Structures to Support Space-Time Dimensional Execution

FIG. 4 illustrates how memory is partitioned between stack and heap in
20 accordance with an embodiment of the present invention. In FIG. 4, memory 400 is divided into a number of regions including heap 402, stacks for threads 404 and speculative heap 406. Heap 402 comprises a region of memory from which objects are allocated. Heap 402 is further divided into younger generation region 408 and older generation region 410 for garbage collection purposes. For performance
25 reasons, garbage collectors typically treat younger generation objects differently from older generation objects. Stack for threads 404 comprises a region of memory from which stacks for various threads are allocated. Speculative heap 406 contains the space-time dimensioned values of all memory elements where the two space-time dimensions of the heap are not collapsed. This includes space-time dimensional
30 versions of objects, for example, version 510 of object 500 as shown in FIG. 5, and objects created by speculative thread 203. For garbage collection purposes, these

objects created by speculative thread 203 can be treated as belonging to a generation that is younger than objects within younger generation region 408.

FIG. 5 illustrates the structure of a primary version of object 500 and a space-time dimensioned version of object 510 in accordance with an embodiment of the present invention.

Primary version of object 500 is referenced by object reference pointer 501. Like any object defined within an object-oriented programming system, primary version of object 500 includes data region 508, which includes one or more fields containing data associated with primary version of object 500. Primary version of object 500 also includes method vector table pointer 506. Method vector table pointer 506 points to a table containing vectors that point to the methods that can be invoked on primary version of object 500.

Primary version of object 500 also includes space-time dimensioned version pointer 502, which points to space-time dimensioned version of object 510, if the two space-time dimensions are not collapsed at this object. Note that in the illustrated embodiment of the present invention, space-time dimensioned version 510 is always referenced indirectly through space-time dimensioned version pointer 502. Primary version of object 500 additionally includes status word 504, which contains status information specifying which fields from data region 508 have been written to or read by speculative thread 203. Space-time dimensioned version of object 510 includes only data region 518.

FIG. 6 illustrates the structure of status word 504 in accordance with an embodiment of the present invention. In this embodiment, status word 504 includes checkpoint number 602 and speculative bits 603. Speculative bits 603 includes read bits 604 and write bits 606. When status word 504 needs to be updated due to a read or a write by speculative thread 203, checkpoint number 602 is updated with the current time of the system. The current time in the time dimension of the system is advanced discretely at a join or a rollback. This allows checkpoint number 602 to be used as a qualifier for speculative bits 603. If checkpoint number 602 is less than the current time, speculative bits 603 can be interpreted as reset.

Read bits 604 keep track of which fields within data region 508 have been read since the last join or rollback. Correspondingly, write bits 606 keep track of which fields within data region 508 have been written since the last join or rollback. In one embodiment of the present invention, read bits 604 includes one bit for each field within data region 508. In another embodiment, read bits includes fewer bits than the number of fields within data region 508. In this embodiment, each bit within read bits 604 corresponds to more than one field in data region 508. For example, if there are eight read bits, each bit corresponds to every eighth field. Write bits 606 similarly can correspond to one or multiple fields within data region 508.

Space-Time Dimensional Update Process

Space-time dimensioning occurs during selected memory updates. For local variable and operand accesses to the system stack, no space-time dimensioned versions exist and nothing special happens. During read operations by head thread 202 to objects in the heap 402, again nothing special happens.

Special operations are involved in write operations by head thread 202 as well as read and write operations by speculative thread 203. These special operations are described in more detail with reference to FIGs. 7, 8 and 9 below.

FIG. 7 is a flow chart illustrating operations involved in a write operation to an object by a head thread 202 in accordance with an embodiment of the present invention. The system writes to the primary version of object 500 and the space-time dimensioned version of object 510 if the two space-time dimensions are not collapsed at this point (step 702). Next, the system checks status word 504 within primary version of object 500 to determine whether a rollback is required (step 704). A rollback is required if speculative thread 203 previously read the data element. The same memory element can be read when the two space-time dimensions of the heap are collapsed at this memory element at the time of the read by speculative thread 203. A rollback is also required if speculative thread 203 previously wrote to the object and thus ensured that the two dimensions of the object are not collapsed at this element, and if the current write operation updates both primary version of object 500 and space-time dimensioned version of object 510.

If a rollback is required, the system causes speculative thread 203 to perform a rollback operation (step 706). This rollback operation allows speculative thread 203 to read from (or write to) the object after head thread 202 writes to the object.

Note that in the embodiment of the present invention illustrated in FIG. 7 the system performs writes to both primary version 500 and space-time dimensioned version 510. In an alternative embodiment, the system first checks to determine if speculative thread 203 previously wrote to space-time dimensioned version 510. If not, the system writes to both primary version 500 and space-time dimensioned version 510. If so, the system only writes to primary version 500.

FIG. 8 is a flow chart illustrating operations involved in a read operation to an object by speculative thread 203 in accordance with an embodiment of the present invention. During this read operation, the system sets a status bit in status word 504 within primary version of object 500 to indicate that primary version 500 has been read (step 802). Speculative thread 203 then reads space-time dimensioned version 510, if it exists. Otherwise, speculative thread 203 reads primary version 500.

FIG. 9 is a flow chart illustrating operations involved in a write operation to a memory element by speculative thread 203 in accordance with an embodiment of the present invention. If a space-time dimensioned version 510 does not exist, the system creates a space-time dimensioned version 510 in speculative heap 406 (step 902). The system also updates status word 504 to indicate that speculative thread 203 has written to the object if such updating is necessary (step 903). The system next writes to space-time dimensioned version 510 (step 904). Such updating is necessary if head thread 202 must subsequently choose between writing to both primary version 500 and space-time dimensioned version 510, or writing only to primary version 500 as is described above with reference to FIG. 7.

FIG. 10 is a flow chart illustrating operations involved in a join operation between head thread 202 and a speculative thread 203 in accordance with an embodiment of the present invention. A join operation occurs for example when head thread 202 reaches a point in the program where speculative thread 203 began executing. The join operation causes state associated with the speculative thread 203 to be merged with state associated with the head thread 202. This involves copying

and/or merging the stack of speculative thread 203 into the stack of head thread 202 (step 1002). It also involves merging space-time dimension and primary versions of objects (step 1004) as well as possibly garbage collecting speculative heap 406 (step 1006). In one embodiment of the present invention, one of threads 202 or 203
5 performs steps 1002 and 1006, while the other thread performs step 1004.

FIG. 11 is a flow chart illustrating operations involved in a join operation between head thread 202 and a speculative thread 203 in accordance with another embodiment of the present invention. In this embodiment, speculative thread 203 carries on as a pseudo-head thread. As a pseudo-head thread, speculative thread 203
10 uses indirection to reference space-time dimensioned versions of objects, but does not mark objects or create versions. While speculative thread 203 is acting as a pseudo-head thread, head thread 202 updates primary versions of objects.

Selectively Marking Objects

FIG. 12 is a flow chart illustrating the process of producing code so that speculative thread 203 does not mark selected memory elements in accordance with an embodiment of the present invention. The system starts by allowing a programmer to examine source code for a program that is to be speculatively executed (step 1202). This enables the programmer to identify read operations within the source code that
20 return values that do not affect subsequent results produced by the computation (step 1204). For example, during a lookup into a hash table, the lookup may involve reading a number of intervening entries in the hash table before finding the desired entry. The values read from these intervening entries do not affect the computation. Hence, the intervening entries do not need to be marked to trigger a subsequent
25 rollback if the intervening entries are subsequently modified by the head thread.

Next, the system allows the programmer to insert hints into the source code telling the compiler to create executable code for speculative thread 203 that does not to mark a given memory element during a given read operation (step 1206). In one embodiment of the present invention, these hints are inserted into the source code in
30 the form of comments. These comments can be read by the compiler, or alternatively by a pre-processor for the compiler.

Next, the system compiles the source code (including the hints) into executable code the causes speculative thread 203 not to mark certain memory elements during specified read operations (step 1208).

Note that this compilation process can take place dynamically during program run time. Alternatively, this compilation process can take place off-line prior to program run time.

FIG. 13 is a flow chart illustrating the process of compiling source code into executable code so that speculative thread 203 does not mark selected memory elements in accordance with an embodiment of the present invention. The system starts by compiling the source code into executable code for both head thread 202 and speculative thread 203 (step 1302).

During this compilation process (or during an associated pre-compilation process), the system reads a comment within the source code specifying that speculative thread 203 is not to mark a given memory element during a given read operation (step 1304). In response to this comment, the system compiles the source code into executable code so that speculative thread 203 does not mark the given memory element during the given read operation (step 1306).

In one embodiment of the present invention this is accomplished by including a "look" attribute for a field within an object indicating that the field is not to be marked by speculative thread 203. At run time, speculative thread 203 examines this look attribute to determine whether a field within an object should be marked during a read operation.

In another embodiment of the present invention, this is accomplished by creating source code for speculative thread 203 that does not mark the given memory element during the given read operation.

Note that by reducing the number of memory elements that are marked, unnecessary rollbacks are eliminated, which thereby improves system performance. Fewer memory elements marked results in fewer loads and stores, and therefore improves system performance.

The foregoing descriptions of embodiments of the invention have been presented for purposes of illustration and description only. They are not intended to

16

be exhaustive or to limit the invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art.

Additionally, the above disclosure is not intended to limit the invention. The scope of the invention is defined by the appended claims.

5

What Is Claimed Is:

1. A method for marking memory elements based upon how information
retrieved from the memory elements is used in order to facilitate speculative program
5 execution, the method comprising:
 examining source code that is to be compiled into executable code for a head
thread that executes program instructions and for a speculative thread that executes
program instructions in advance of the head thread;
 wherein during read operations to memory elements by the speculative thread,
10 the executable code generally causes the speculative thread to update status
information associated with the memory elements to indicate that the memory
elements have been read by the speculative thread;
 identifying a given read operation from a given memory element, wherein a
given value retrieved from the given memory element during the given read operation
15 is not used in by the speculative thread;
 inserting at least one hint into the source code specifying that the speculative
thread is not to update status information during the given read operation directed to
the given memory element; and
 compiling the source code, including the at least one hint, into the executable code, so
20 that during the given read operation, the executable code does not cause the
speculative thread to update status information associated with the given memory
element to indicate that the given memory element has been read by the speculative
thread.
- 25 2. The method of claim 1, wherein during a subsequent write operation to a
memory element by the head thread, the executable code causes the head thread to:
 check status information associated with the memory element to determine if
the memory element has been read by the speculative thread; and
 if the memory element has been read by the speculative thread, causes the
30 speculative thread to roll back in order to read a result of the write operation.

18

3. The method of claim 1, wherein inserting the at least one hint into the source code involves inserting a comment into the source code that causes a compiler to generate executable code for the speculative thread that does not to mark the given memory element during the given read operation.

5

4. The method of claim 1, wherein compiling the source code, including the at least one hint, into the executable code, involves setting a look attribute associated with the given memory element to indicate that the speculative thread is not to mark the given memory element.

10

5. The method of claim 1, wherein the executable code causes a join operation between the head thread and the speculative thread when the head thread reaches a point in a program where the speculative thread began executing, the join operation causing state associated with the speculative thread to be merged with state associated with the head thread.

15

6. The method of claim 1, further comprising integrating the source code into a system code library.

20

7. The method of claim 1, wherein the memory elements include objects defined within an object-oriented programming system.

25

8. A method for compiling source code into executable code for a head thread that executes program instructions and a speculative thread that executes program instructions in advance of the head thread, the method comprising:

compiling the source code into executable code so that during read operations to memory elements by the speculative thread, the executable code generally causes the speculative thread to update status information associated with the memory elements to indicate that the memory elements have been read by the speculative thread;

30

during the compilation process, reading at least one hint inserted into the source code specifying that the speculative thread is not to update status information during a given read operation directed to a given memory element; and

in response to the at least one hint, compiling the source code into the
5 executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

10 9. The method of claim 8, wherein during a subsequent write operation to a memory element by the head thread, the executable code causes the head thread to:
check status information associated with the memory element to determine if the memory element has been read by the speculative thread; and
if the memory element has been read by the speculative thread, causes the
15 speculative thread to roll back in order to read a result of the write operation.

10. The method of claim 8, wherein reading the at least one hint includes reading a comment within the source code.

20 11. The method of claim 8, wherein compiling the source code includes compiling the source code dynamically as the executable code is being executed.

12. A apparatus for compiling source code into executable code for a head thread that executes program instructions and a speculative thread that executes
25 program instructions in advance of the head thread, comprising:

a compiler that is configured to compile source code into executable code so that during read operations to memory elements by the speculative thread, the executable code generally causes the speculative thread to update status information associated with the memory elements to indicate that the memory elements have been
30 read by the speculative thread;

wherein the compiler is additionally configured to,

20

read at least one hint inserted into the source code specifying that the speculative thread is not to update status information during a given read operation directed to a given memory element, and to

5 compile the source code into the executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

10 13. The apparatus of claim 12, wherein during a subsequent write operation to a memory element by the head thread, the executable code causes the head thread to:

check status information associated with the memory element to determine if the memory element has been read by the speculative thread; and

15 if the memory element has been read by the speculative thread, causes the speculative thread to roll back in order to read a result of the write operation.

20 14. The apparatus of claim 12, wherein the compiler is configured to read the at least one hint in the form of a comment within the source code.

15 15. The apparatus of claim 12, wherein the compiler is configured to set a look attribute associated with the given memory element to indicate that the speculative thread is not to mark the given memory element.

25 16. The apparatus of claim 12, wherein the compiler is configured to produce the executable code so that the executable code causes a join operation between the head thread and the speculative thread when the head thread reaches a point in a program where the speculative thread began executing, the join operation causing state associated with the speculative thread to be merged with state associated
30 with the head thread.

21

17. The apparatus of claim 12, wherein the memory elements include objects defined within an object-oriented programming system.

18. The apparatus of claim 12, wherein the compiler is configured to
5 compile the source code dynamically as the executable code is being executed.

19. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform a method for compiling source code into executable code for a head thread that executes program instructions and a
10 speculative thread that executes program instructions in advance of the head thread, the method comprising:

compiling the source code into executable code so that during read operations to memory elements by the speculative thread, the executable code generally causes the speculative thread to update status information associated with the memory
15 elements to indicate that the memory elements have been read by the speculative thread;

during the compilation process, reading at least one hint inserted into the source code specifying that the speculative thread is not to update status information during a given read operation directed to a given memory element; and

20 in response to the at least one hint, compiling the source code into the executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

25

20. The computer-readable storage medium of claim 19, wherein during a subsequent write operation to a memory element by the head thread, the executable code causes the head thread to:

check status information associated with the memory element to determine if
30 the memory element has been read by the speculative thread; and

22

if the memory element has been read by the speculative thread, causes the speculative thread to roll back in order to read a result of the write operation.

21. The computer-readable storage medium of claim 19, wherein reading
5 the at least one hint includes reading a comment within the source code.

22. The computer-readable storage medium of claim 19, wherein
compiling the source code includes compiling the source code dynamically as the
executable code is being executed.

10

23. The computer-readable storage medium of claim 19, wherein
compiling the source code, including the at least one hint, into the executable code,
involves setting a look attribute associated with the given memory element to indicate
that the speculative thread is not to mark the given memory element.

15

24. The computer-readable storage medium of claim 19, wherein the
executable code causes a join operation between the head thread and the speculative
thread when the head thread reaches a point in a program where the speculative thread
began executing, the join operation causing state associated with the speculative
20 thread to be merged with state associated with the head thread.




FIG. 1




FIG. 2A
(PRIOR ART)




FIG. 2B




FIG. 3




FIG. 4

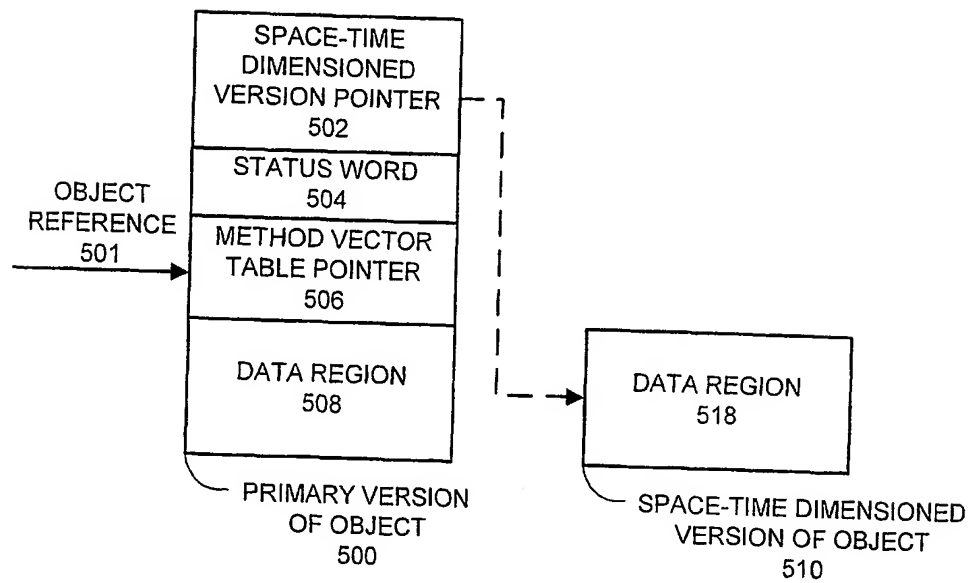


FIG. 5




FIG. 6

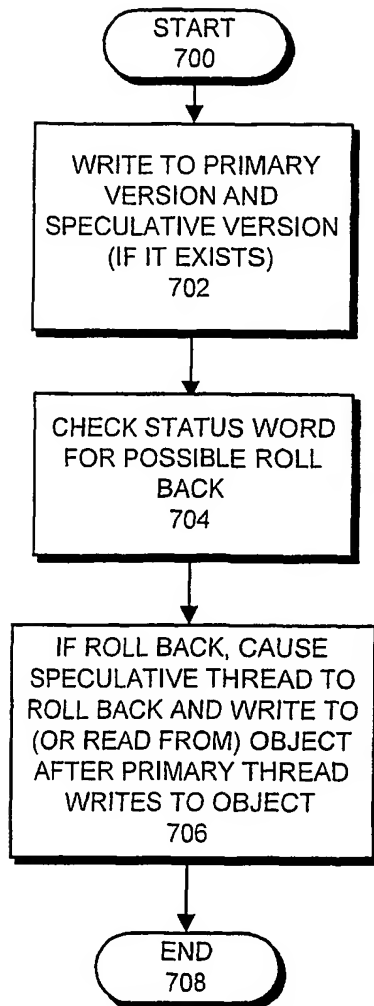


FIG. 7




FIG. 8




FIG. 9




FIG. 10


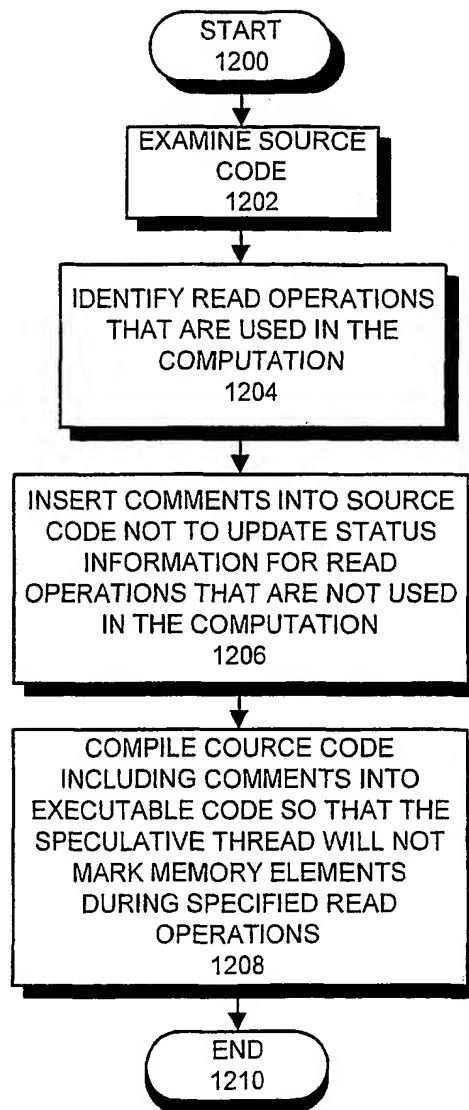


FIG. 11

**FIG. 12**

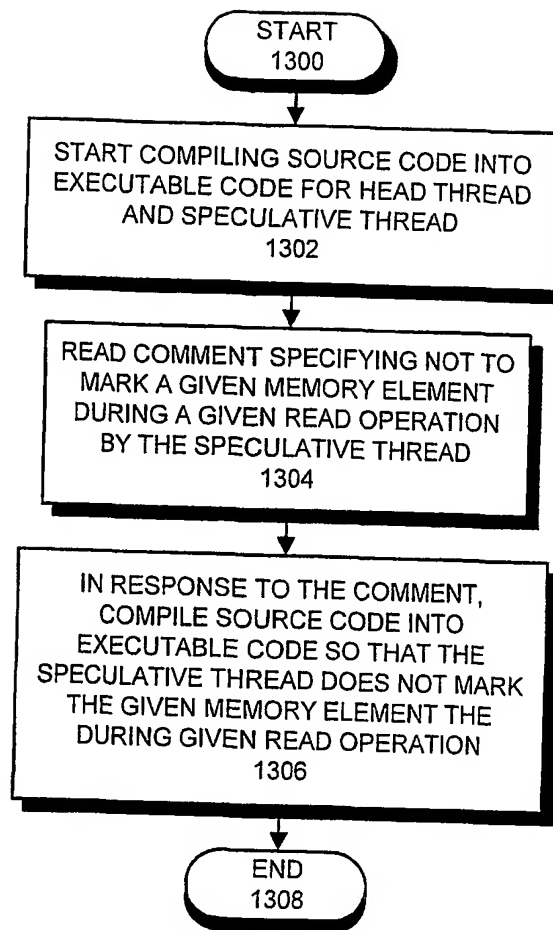


FIG. 13

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 December 2001 (06.12.2001)

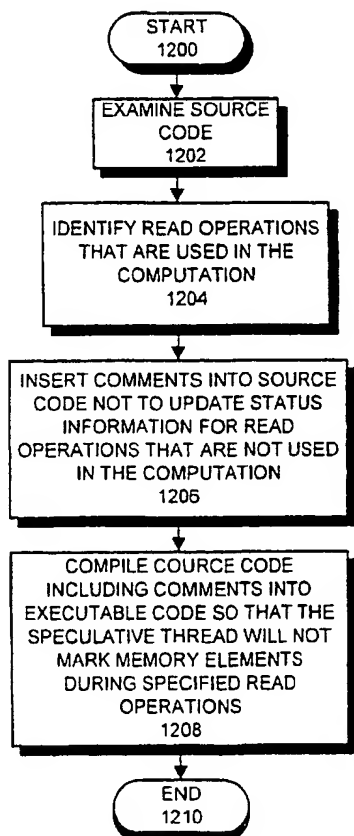
PCT

(10) International Publication Number
WO 01/93028 A3

- (51) International Patent Classification⁷: **G06F 9/38**
- (21) International Application Number: **PCT/US01/17645**
- (22) International Filing Date: **30 May 2001 (30.05.2001)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
60/208,428 31 May 2000 (31.05.2000) US
09/761,226 16 January 2001 (16.01.2001) US
- (71) Applicant: **SUN MICROSYSTEMS, INC.** [US/US]; 901 San Antonio Road, Palo Alto, CA 94303 (US).
- (72) Inventors: **CHAUDHRY, Shailender**; 1200 Gough St., Apt. 10F, San Francisco, CA 94109 (US). **TREMBLAY, Marc**; 140 Hanna Way, Menlo Park, CA 94025 (US).
- (74) Agent: **PARK, Richard**; 508 2nd Street, Suite 201, Davis, CA 95616 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report

[Continued on next page]


(54) Title: MARKING MEMORY ELEMENTS BASED UPON ACCESS INFORMATION DURING SPECULATIVE EXECUTION



(57) Abstract: One embodiment of the present invention provides a system that marks memory elements based upon how information retrieved from the memory elements affects speculative program execution. This system operates by allowing a programmer to examine source code that is to be compiled into executable code for a head thread that executes program instructions, and for a speculative thread that executes program instructions in advance of the head thread. During read operations to memory elements by the speculative thread, this executable code generally causes the speculative thread to update status information associated with the memory elements to indicate that the memory elements have been read by the speculative thread. Next, the system allows the programmer to identify a given read operation directed to a given memory element, wherein a given value retrieved from the given memory element during the given read operation does not affect subsequent execution of the speculative thread. The programmer is then allowed to insert a hint into the source code specifying that the speculative thread is not to update status information during the given read operation directed to the given memory element. Next, the system compiles the source code, including the hint, into the executable code, so that during the given read operation, the executable code does not cause the speculative thread to update status information associated with the given memory element to indicate that the given memory element has been read by the speculative thread.

WO 01/93028 A3

WO 01/93028 A3



(88) Date of publication of the international search report:
4 April 2002

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 01/17645

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/38

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P, A	WO 00 70451 A (SUN MICROSYSTEMS INC) 23 November 2000 (2000-11-23) the whole document ---	1-24
A	EP 0 725 334 A (IBM) 7 August 1996 (1996-08-07) abstract page 5, line 32 - line 48 page 6, line 33 - line 42 page 7, line 49 - line 53 page 8, line 51 - page 9, line 2 page 10, line 53 - page 11, line 8 page 11, line 45 - line 54 page 12, line 21 - line 39 page 13, line 6 - line 15 --- -/--	1-24

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

Z document member of the same patent family

Date of the actual completion of the international search

12 December 2001

Date of mailing of the international search report

20/12/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel (+31-70) 340-2040 Tx 31 651 epo nl
Fax (+31-70) 340-3016

Authorized officer

Moraiti, M

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 01/17645

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>"CHARACTER STRINGE" . LECTURE NOTES IN COMPUTER SCIENCE. SPRINGER VERLAG, NEW YORK, NY, US, VOL. 106, PAGE(S) 11-12, 86-87, 127, 139-140, 211-212 XP002026321 ISSN: 0302-9743 page 12</p>	1, 8, 12, 19
A	<p>GB 2 321 544 A (IBM) 29 July 1998 (1998-07-29)</p>	

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No

PLI/US 01/17645

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0070451	A	23-11-2000	WO 0070450 A1	23-11-2000
			WO 0070451 A1	23-11-2000
			WO 0070452 A1	23-11-2000
EP 0725334	A	07-08-1996	US 5812811 A	22-09-1998
			EP 0725334 A1	07-08-1996
			JP 3093626 B2	03-10-2000
			JP 8249183 A	27-09-1996
GB 2321544	A	29-07-1998	NONE	

